

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



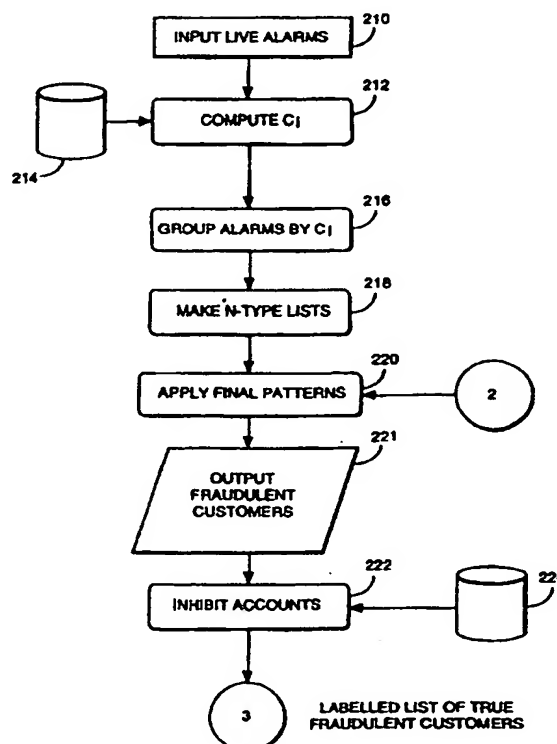
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H04M 15/00, 3/38, H04Q 3/00		A1	(11) International Publication Number: WO 97/37486
			(43) International Publication Date: 9 October 1997 (09.10.97)
(21) International Application Number: PCT/GB97/00836			(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ARIPO patent (GH, KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).
(22) International Filing Date: 25 March 1997 (25.03.97)			
(30) Priority Data: 96302240.5 29 March 1996 (29.03.96) EP			
(34) Countries for which the regional or international application was filed: GB et al.			
(71) Applicant (for all designated States except US): BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81 Newgate Street, London EC1A 7AJ (GB).			
(72) Inventor; and (75) Inventor/Applicant (for US only): BUSUIOC, Nicolae, Marius [RO/GB]; 28 Henslow Road, Ipswich, Suffolk IP4 5EG (GB).			
(74) Agent: BRADLEY, David, William; BT Group Legal Services, Intellectual Property Dept., 8th floor, 120 Holborn, London EC1N 2TE (GB).			Published With international search report.

(54) Title: **FRAUD MONITORING IN A TELECOMMUNICATIONS NETWORK**

(57) Abstract

A method of and system for detecting the possible fraudulent use of a telecommunications network involves applying rule-based criteria to generate a plurality of fraud alarms, each corresponding to an individual rule. Each alarm is associated with a particular customer, and for each individual customer a note is made of the total alarms generated by that customer and the grouping of individual alarm types generated. The customer's call is then determined to be fraudulent or otherwise based upon prior experience of past customers who have generated that particular profile of alarm grouping and total number of alarms. The system automatically outputs a list of potentially fraudulent customers, the accounts of which may either be further investigated or may automatically be inhibited.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

FRAUD MONITORING IN A TELECOMMUNICATIONS NETWORK

The present invention relates a telecommunications network and more particularly to a method of, and a system for, detecting the possible fraudulent use
5 of a telecommunications network.

Rule-based fraud detection systems attempt to detect fraudulent usage by comparing details of individual calls over the telecommunications network with a series of one or more predefined rules. If a particular usage of the network (to be referred to throughout this specification as a "call record") triggers one or more of
10 the predefined rules, an alarm is generated, enabling human operators to take the necessary action. While such systems have had some success in combating fraud, difficulties tend to arise due to the sheer number of alarms that can be generated within a short time. Typically, fraud detection operators may have tens of thousands of live alarms to deal with during a day, and it is therefore generally
15 impractical to deal with each individual alarm as it arises. Methods have been developed for consolidating or grouping the fraud alarms based on their priority, but the workload for the fraud operators still remains substantial.

Work has been done to provide correlated fault alarms for identifying possible faulty network devices and/or failure of communication links in
20 telecommunication networks. However, the correlation process here relies very much upon the fact that the network topology is well known, with the alarms and the alarms correlations being calculated on that basis.

It is an object of the present invention at least to alleviate these problems. It is a further object to provide a method of, and a system for, detecting the
25 possible fraudulent use of a telecommunications network which can be used across a range of products and services.

According to a first aspect of the present invention there is provided a method of detecting the possible fraudulent use of a telecommunications network, the method comprising:

- 30 (a) receiving alarms indicative of potentially fraudulent calls on the network, the alarms being divided into a plurality of alarm types;
- (b) associating a unique customer identifier with each alarm;

(c) selecting a test class of customer identifiers such that each customer identifier in the test class is associated with a given grouping of alarm types;

(d) identifying those customer identifiers within the test class that are
5 associated with known fraudulent calls and deriving a measure therefrom indicative of fraud within the test class; and

(e) determining that any customer identifier associated with further alarms is connected with fraudulent use of the network if it falls within the said test class and if the said measure for that class exceeds a given level.

10 According to a second aspect of the invention there is provided a system for detecting the possible fraudulent use of a telecommunications network, the system comprising:

(a) means for receiving alarms indicative of potentially fraudulent calls on the network, the alarms being divided into a plurality of alarm types;

15 (b) means for associating a unique customer identifier with each alarm;

(c) means for selecting a test class of customer identifiers such that each customer identifier in the test class is associated with a given grouping of alarm types;

(d) means for identifying those customer identifiers within the test class
20 that are associated with known fraudulent calls and deriving a measure therefrom indicative of fraud within the test class; and

(e) means for determining that any customer identifier associated with further alarms is connected with fraudulent use of the network if it falls within the said test class and if the said measure for that class exceeds a given level.

25 By iterating the method, the system gradually learns and becomes more effective at identifying fraud.

The invention discovers patterns in the alarm data, and operates on those, rather than operating on the rules that generate the alarms themselves. Preferably, the system attempts to detect fraudulent usage by measuring and
30 comparing the parameters values of individual calls, over the telecommunications network, against pre-set thresholds within the detection rules. This allows for a reduced number of derived alarms to be created, thereby easing the task of the fraud operators. In contrast with known network fault alarm correlations, the

invention is not limited to use on any specific network or on any specific model. Instead, it identifies fraud trends by identifying patterns in particular groupings of raw alarms. The solution is applicable across all products and services.

In one form, the invention may provide the fraud operators with a display
5 identifying, in order, those groups of alarms which are most indicative of the presence of fraud and, against each group, a list of (normalized) customer identifiers whose calls have triggered alarms in that particular group. A numerical measure may be associated with each grouping, providing the fraud operators with a quantitative estimate of the probability that a particular customer identifier is
10 associated with fraudulent calls.

The system may automatically determine that certain alarm groupings are associated with fraud (for example if the measure exceeds a predefined value), and may automatically inhibit the user accounts corresponding to the user identifiers which fall within those groupings. Alternatively, the information may be provided
15 to human operators, who may reserve to themselves the final decisions.

It is not essential, of course, that the measure takes the form of a single numerical value. It could, instead, consist of several numerical or non-numerical indicators that may be tested against a predefined level. Again, the given level in that case need not itself be a single numerical value. It will be understood, of
20 course, that if the measure increases with fraud, then it will exceed the given level in the upward-going direction when the measure is larger than the level. On the other hand, if the measure is designed to fall with increasing fraud, then it will exceed the given level in the downward-going direction when it falls to a value below that of the given level.

25 In its various forms, the invention, or preferred aspects of it, may provide a very concise easily-understood presentation of alarm information to the fraud operator. It provides improved use of alarm data, along with the flexibility to add new alarm types and continuously to detect and learn new alarm types. It allows easier detection of fraud by the human operator, or alternatively may be arranged
30 to detect fraud automatically. This may provide substantial revenue savings from the increased ability of the fraud detection systems, as a whole, to detect fraud at an early stage and to apply preventative measures.

The invention may be carried into practice in a number of ways and one specific embodiment will now be described, by way of example, with reference to the accompanying figures, in which:

Figure 1 shows how the system is trained and the alarms patterns are
5 refined over time; and

Figure 2 shows how the system is used on real data but continues to learn new patterns through performance evaluation.

The fraud detection method and system shown in Figures 1 and 2 may typically be embodied in a computer program running on a dedicated server which
10 is attached to the telecommunications network to be monitored. Depending on the size of the network, there may be a single server, or the system may be duplicated on several servers, spaced across the network. All or parts (modules) of the system could alternatively be hard-coded rather than being embodied by way of a computer program, especially the modules engaged in pure computation. The
15 system is designed to receive information from external sources across the network, in the form of a plurality of fraud alarms A_i . These alarms are generated by testing each call that is made on the telecommunications network against a corresponding rule set, with the alarm being automatically activated if the call matches the requirements of the rule. The rules are preferably independent, or at
20 least partially so, so that if for example a single call activates alarms A_1 and A_2 , the existence of both alarms provides some additional evidence by way of cross-check that the call is indeed fraudulent, over and above the information that would be provided by one of the alarms alone. One rule might state, for example, that fraud is a possibility if the call is an international call being made from a public call
25 box to a country known to be a supplier of illegal drugs. Another rule might suggest fraud if the call has been paid for by charge-card, and the call does not fit the call history on that particular account. A further rule might suggest that fraud is taking place if a low-usage charge-card customer suddenly starts making a long series of international telephone calls to different countries from a public 'phone
30 box.

Each alarm A_i may be associated with a particular customer C_i who is paying the bill for the call that generated that alarm.

For ease of description the preferred embodiment will be described with reference to Figures 1 and 2, and concurrently in association with a hypothetical worked example. The example will assume that alarms are generated by four different rules, giving rise to four separate alarm types A_1 , A_2 , A_3 and A_4 . It will also be assumed that the network has nine customers, identified respectively as C_1 to C_9 .

Before the system may be operated on live data, it first has to be trained through various training cycles using a set of pre-existing alarm data. This is data that has been already analysed by the fraud operators who have labelled each alarm accordingly as indicative of real fraud or not fraud. Turning first to Figure 1, the test alarms are received by the system at 10, and the corresponding C_i for each alarm A_i is then determined at 12. To assist in this determination, information from an external or internal database 14 may be used. If the customer is a direct customer of the telecommunications network owner, customer details may be looked up directly in the corresponding customer database. On the other hand, the customer may have connected into the network being monitored via another network (perhaps abroad) and the fraud operators may accordingly have no direct access to customer and billing details for that particular customer. In such a case, the customer is merely identified for the purposes of the system by a unique reference number; this could come from information provided by the owner of the foreign network or, in the absence of that, from a combination of the calling number and the called number.

Once each alarm A_i has been associated with a particular customer C_i , the information is passed on to a statistical analyser module 15 which first groups the alarms by C_i as indicated at 16. A typical grouping for the example being considered might be as follows:

Training Set:

C_i	A				X_i
	1	2	3	4	
C_1	1	1	1	2	5
C_2	0	0	3	2	5
C_3	1	1	1	1	4
C_4	1	1	1	1	4
C_5	2	1	1	1	5
C_6	1	2	0	0	3
C_7	0	0	3	2	5
C_8	1	2	1	1	5
C_9	0	0	2	3	5

Each cell in the table represents the number of times a particular alarm has been triggered for that particular customer. For example, customer C_1 has triggered alarms A_1 , A_2 and A_3 once each, and has triggered alarm A_4 twice. The last column in the table, labelled X_i , simply represents the total number of alarms of all types for customer C_i .

The alarms in the training set are now re-grouped into the "N-type list" shown below. This is a table in which each row represents one of the possible groupings G of the alarms A_i , as determined from the training set. Each column of the table represents the total number of alarms of all types, X_i . Each customer C_i appears exactly once in the table.

N-type List:

G	5	4	3
$A_1 A_2 A_3 A_4$	$C_1 C_5 C_8$	$C_3 C_4$	
$A_1 A_2$			C_6
$A_3 A_4$	$C_2 C_7 C_9$		

It can be seen from the table that there are three customers who actuated a total of five alarms in all four alarm types, namely C_1 , C_5 and C_8 .

The information from the N-type list is passed to a pattern extractor module 19, which first analyses it, at 20. Input is provided from the training set 22 and from an external database 24 which might include, for example, details of the billing and/or call histories of the customers being investigated. The analysis is carried out by asking skilled operators to check the training set, and to determine for each of the customers C_i whether that customer is in fact fraudulent or not. From the information in the external database, the fraud operators may be able to say with some certainty that particular customers are indeed acting fraudulently. For this particular example it will be assumed that customers C_1 , C_3 , C_5 and C_9 are fraudulent; these are shown in bold type in the N-type list above.

At step 26, alarm patterns are produced by considering individually each populated cell within the table. Taking first the top left cell, it will be apparent that two of the three customers in that cell are considered fraudulent, so it can be said that the chance of a customer falling into that cell being fraudulent is some 67%. In the adjacent cell, containing C_3 and C_4 , only one of the two customers has been flagged as fraudulent, so the ratio for that cell is 50%. Continuing with the other cells, one can produce an alarm pattern table as follows:

F_i	X_i	G	C_i
67%	5	$A_1 A_2 A_3 A_4$	$C_1 C_5 C_8$
50%	4	$A_1 A_2 A_3 A_4$	$C_3 C_4$
33%	5	$A_3 A_4$	$C_2 C_7 C_9$
0%	3	$A_1 A_2$	C_6

In the alarm pattern table, each populated cell in the N-type list has its own row, the rows being ordered by the value of F, the likelihood of a customer who appears in that row being fraudulent.

In the above table, the group $A_1 A_2 A_3 A_4$ appears twice, and the table is now refined so that each alarm group is represented by a single row. The combined likelihood for the first two rows can be computed using the formula:

$$F = \sum_j (X_{ij} F_{ij}) / \sum_j X_{ij}$$

where j = number of distinct X_i for this alarm group

(here $j = 2$, since $X_i = 4$ and 5)

F_{ij} = the partial likelihoods of each row (here 67% and 50%).

This gives a combined likelihood for the first two rows of

$$F = (5 \times 0.67 + 4 \times 0.50) / (5 + 4) = 59.4\%.$$

This, then, is the overall likelihood of fraud within the group $A_1 A_2 A_3 A_4$.

Each combined or individual likelihood F_i may have associated with it a
 5 confidence value K_i (not shown in the table above). This may be computed as the
 ratio of fraudulent customers in this group to the number of fraudulent customers
 in all groups detected in the current training set.

After rebuilding the table and ordering by combined likelihood values, one
 may see from the C_i column that customers C_1 , C_5 and C_3 are fraudulent, and that
 10 C_4 and C_8 are suspicious. The system may automatically keep track of the
 fraudulent and suspicious customers by storing them in appropriate databases,
 including a suspicious customer database 30 (Figure 1).

The current table now reads as follows:

F_i	K_i	X_i	G	C_i
59.4%	75%	5,4	$A_1 A_2 A_3 A_4$	$C_1 C_5 C_8 C_3 C_4$
33%	25%	5	$A_3 A_4$	$C_2 C_7 C_9$
0%		3	$A_1 A_2$	C_6

15 Where the values K_i are here assumed to have been calculated as shown based
 upon the rest of the data set.

It will be noted that customers C_8 and C_4 appear in the top row, along
 with C_1 and C_5 , indicating that they may be fraudulent as well. Customer C_9 , on
 the other hand, appears in the second row with C_2 and C_7 ; so C_2 and C_7 have to
 20 be treated as suspicious.

The system now learns the patterns set out in the alarm pattern table, at
 step 28 of Figure 1. The learned patterns include a cut-off point above which a
 customers are to be deemed potentially fraudulent; here, the cut-off point may be
 for example 30%, so that groups $A_1 A_2 A_3 A_4$ and $A_3 A_4$ may be considered as
 25 indicative of at least potential fraud.

In an alternative embodiment, the value of F may be calculated in some
 more sophisticated way simply than taking the number of known fraudulent
 customers in a particular cell, and dividing by the total number of customers in that

cell. The figures might be weighted, for example, in dependence upon the cost of the potential fraud. This could be derived from the cost of all of the individual calls made by the customer which have produced alarms, or all of such calls that have taken place over a defined period such as the last seven days. Other criteria
5 may no doubt occur to the skilled man.

The next stage in the procedure is to refine and update the learned patterns through a new training cycle. A new training data set is provided at the input module 10 and after the same computations as previously described, new alarm types are produced and the old ones are updated, as described below, at 26.

10 Let us assume now that during the second training cycle the group $A_1 A_2 A_3 A_4$ consists of only one fraudulent customer, and that X_i for that customer equals 3. Accordingly, using the same analysis as before, the value of F_i for that group will be 100%. Let us assume, further, that the corresponding confidence K_i equals 10%.

15 The likelihood of fraud for this new alarm group is now updated at 26 using the equation:

$$F_{\text{update}} = K_{\text{old}}F_{\text{old}} + K_{\text{new}}F_{\text{new}}$$

For the present example, this gives:

$$F_{\text{update}} = 59.4 \times 0.75 + 0.1 \times 100$$

20
$$= 54.5\%.$$

Once the patterns have been updated at 26, they are then learned at 28. Once the underlying patterns have been revised as necessary to provide optimal performance, the final alarm pattern is output and/or stored for use on live data.

Turning now to Figure 2, we will describe how the system is run on real
25 alarm data (that is live, unlabelled alarms). The live data to be monitored arrives by way of the live alarms 210, against which are computed the corresponding customers C_i at step 212. As before, an external database 214 may be used as necessary. The alarms are grouped by C_i at 216, and N-type lists constructed at 218.

30 The final alarm pattern table is applied against the N-type list at 220, and any customer who appears in the top row of that list (or more generally, who appears in any of the uppermost rows in which the value F is greater than a defined threshold value) is output at 221 as a fraudulent customer.

The fraudulent customers list 221 is considered by the fraud operators at 222, and those customers who are considered to be truly fraudulent have their accounts inhibited at that stage. In deciding which customers are truly fraudulent, the fraud operators may have access to additional database information, as
5 indicated at 224.

The labelled list of true fraudulent customers is then sent back, as indicated in Figures 1 and 2, to the pattern extractor module 19 where the pattern likelihoods are adjusted accordingly, and new alarm groups are added as necessary.

10 Then the whole process restarted on receipt of a new group of real alarm data at 210 for processing. The process shown in Figure 2 is continually reiterated, with the result that the grouped alarms, the N-type list and the alarm pattern table continually changes according to the customers currently involved and the particular alarms and alarm groups they have generated. The alarm
15 pattern table may be shown on the fraud operator's screen, and will constantly be updated as the groupings and the customers change. As customer accounts are inhibited, customers in the uppermost rows which are now defined as fraudulent continually disappear, with others coming in all the time.

The fraud operators are provided with a readily-comprehensible list of
20 potentially fraudulent customers, (from the suspicious customers database 30), ordered according to the likelihood of fraud. It is therefore relatively easy for the operators to inhibit accounts as necessary, either manually or automatically. A combination of both may be used, for example all customers having a value of F greater than 95% may automatically have their accounts inhibited, and all
25 customers having a value F between 85% and 95% may be considered for further manual investigation.

Some customers may of course not be customers of the telecommunications network which is being monitored, in which case it may not be possible to inhibit their accounts. However, since each customer has a unique
30 reference identifier C_i , the necessary information can be passed to the owners of the external network from which the call is emanating, suggesting that they might investigate this particular customer account.

Manual or automatic investigations may also be made as to the connections between the fraudulent customers, to check for evidence of organised crime.

The threshold in F for determining whether a customer is fraudulent may
5 be varied either manually or automatically as desired. Continually varying the cut-off points avoids the problem of fraudsters getting to know what the cut-off points are, and altering their behaviour accordingly.

It will be understood of course that in a practical system there may be an extremely large number of alarm categories A_i , and a consequently large number of
10 category groups G in the N-type list. There will also be a large number of customers, with the result that the statistical analysis involved in creating the alarm pattern table will be substantially more reliable than may have appeared from the simplistic example that has been used for the purposes of discussion.

In one preferred embodiment, the system may keep a running total over
15 time of the percentage of customers falling into each cell of the N-type list who either automatically or manually have their accounts inhibited as being used fraudulently. This information may be used to provide constantly updated values of F for each cell or alarm grouping, thereby allowing the order of entries in the alarm pattern table to change over time as the fraudster's behaviour varies.

CLAIMS:

1. A method of detecting the possible fraudulent use of a telecommunications network, the method comprising:
 - 5 (a) receiving alarms indicative of potentially fraudulent calls on the network, the alarms being divided into a plurality of alarm types;
 - (b) associating a unique customer identifier with each alarm;
 - (c) selecting a test class of customer identifiers such that each customer identifier in the test class is associated with a given grouping of alarm
 - 10 types;
 - (d) identifying those customer identifiers within the test class that are associated with known fraudulent calls and deriving a measure therefrom indicative of fraud within the test class; and
 - (e) determining that any customer identifier associated with further
 - 15 alarms is connected with fraudulent use of the network if it falls within the said test class and if the said measure for that class exceeds a given level.
2. A method as claimed in Claim 1 in which the measure is a weighted or unweighted function of the number of:
 - 20 (A) customer identifiers within the test class that are associated with known fraudulent calls; and
 - (B) the total number of customer identifiers in the test class.
3. A method as claimed in Claim 2 in which the function is the ratio (A)/(B).
- 25 4. A method as claimed in any one of Claims 1 to 3 in which the measure is a function of the potential costs of the known fraudulent calls related to customer identifiers falling into the test class.
- 30 5. A method as claimed in any one of the preceding claims in which the said given level is user-defined.

6. A method as claimed in any one of the preceding claims in which the said given grouping of alarm types is at least partly defined by a unique combination of available alarm types or of any subset thereof.

5 7. A method as claimed in any one of Claims 1 to 5 in which the said given grouping of alarm types is at least partly defined by:

(a) a unique combination of available alarm types or of any subset thereof; and

(b) the total number of alarms of all types for that combination.

10

8. A method as claimed in any one of the preceding claims including selecting a plurality of test classes associated with a corresponding plurality of given groupings of alarm types.

15 9. A method as claimed in Claim 8 including selecting all possible groupings of alarm types from all unique combinations of available alarm types or of any subset thereof.

10. A method as claimed in Claim 8 or Claim 9 including deriving an individual
20 measure from each of the test classes, and sorting the test classes in order according to the values of the individual measures.

11. A method as claimed in Claim 10 including displaying the test classes in the said order along with information on the customer identifiers falling into each
25 test class.

12. A method as claimed in any one of the preceding claims including the step of inhibiting a user account associated with a customer identifier determined as being connection with fraudulent use of the network.

30

13. A method as claimed in any one of the preceding claims including updating the measure in step (d) on the basis of an independent analysis as to whether the

Zcustomer identifier determined at step (c) to be associated with fraud has been correctly so determined.

14. A method as claimed in Claim 13 when dependent upon Claim 11 including automatically updating the display of test classes.

15. A method as claimed in any one of the preceding claims including maintaining a database of customer identifiers and, associated with the customer identifiers, the number of alarms generated by that customer broken down by alarm type.

16. A method as claimed in Claim 15 in which the database further includes the total number of alarms of all types corresponding to each customer identifier.

17. A system for detecting the possible fraudulent use of a telecommunications network, the system comprising:

(a) means for receiving alarms indicative of potentially fraudulent calls on the network, the alarms being divided into a plurality of alarm types;

(b) means for associating a unique customer identifier with each alarm;

(c) means for selecting a test class of customer identifiers such that each customer identifier in the test class is associated with a given grouping of alarm types;

(d) means for identifying those customer identifiers within the test class that are associated with known fraudulent calls and deriving a measure therefrom indicative of fraud within the test class; and

(e) means for determining that any customer identifier associated with further alarms is connected with fraudulent use of the network if it falls within the said test class and if the said measure for that class exceeds a given level.

Fig.1.


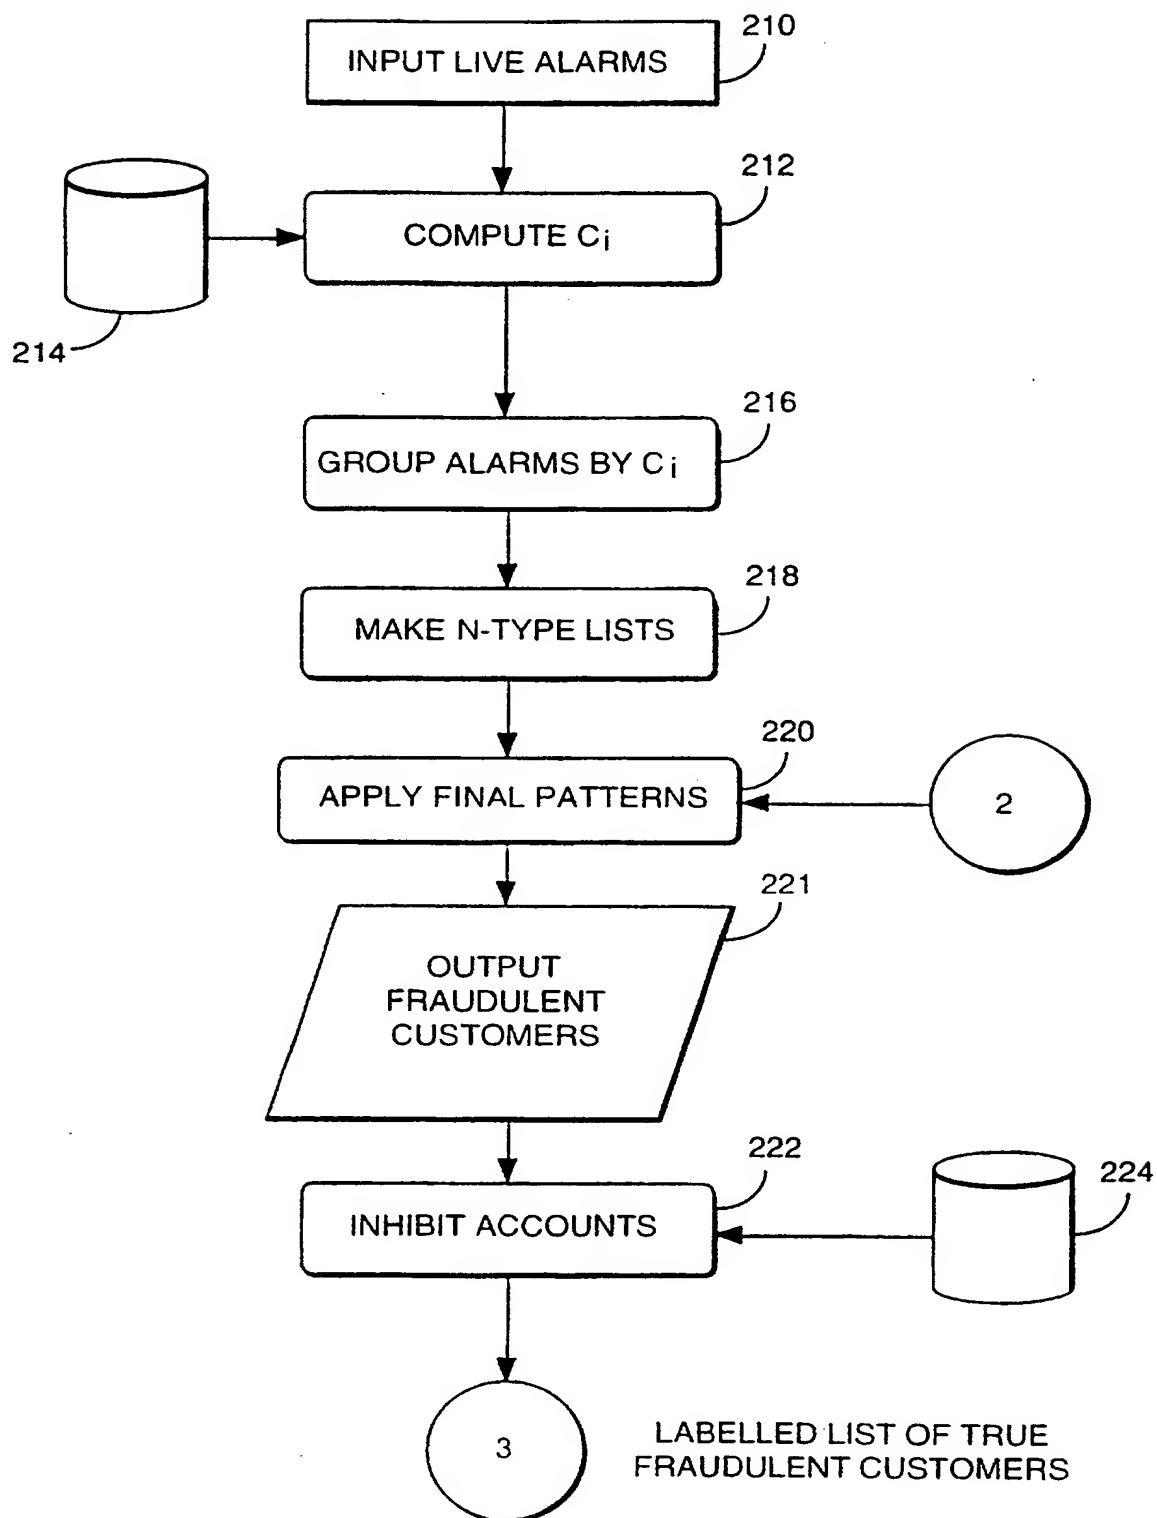


Fig.2.



INTERNATIONAL SEARCH REPORT

Int: nal Application No
PCT/GB 97/00836

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 H04M15/00 H04M3/38 H04Q3/00

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 H04M H04Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	EP 0 653 868 A (AT&T CORP.) 17 May 1995 see the whole document ---	1-17
Y	EP 0 583 135 A (AMERICAN TELEPHONE AND TELEGRAPH COMPANY) 16 February 1994 see the whole document ---	1-17
A	EP 0 618 713 A (AT&T CORP.) 5 October 1994 see column 2, line 16 - line 18 ---	1-17
A	EP 0 661 863 A (AT&T CORP.) 5 July 1995 see the whole document ---	1-17
A	WO 94 11959 A (CORAL SYSTEMS, INC.) 26 May 1994 see abstract ---	1-17
	-/--	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *Z* document member of the same patent family

Date of the actual completion of the international search

29 May 1997

Date of mailing of the international search report

09.06.97

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+ 31-70) 340-3016

Authorized officer

Montalbano, F

INTERNATIONAL SEARCH REPORT

International Application No

PCT/GB 97/00836

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,X	<p>US 5 602 906 A (PHELPS J.W.) 11 February 1997 see the whole document -----</p>	1,17

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/GB 97/00836

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 653868 A	17-05-95	US 5495521 A CA 2132557 A CN 1110042 A JP 7212502 A	27-02-96 13-05-95 11-10-95 11-08-95
EP 583135 A	16-02-94	US 5357564 A CA 2100846 A,C CN 1083295 A JP 7321919 A	18-10-94 13-02-94 02-03-94 08-12-95
EP 618713 A	05-10-94	AU 5767194 A CA 2114155 A JP 6350698 A	06-10-94 01-10-94 22-12-94
EP 661863 A	05-07-95	US 5463681 A CA 2138420 A JP 7212460 A	31-10-95 30-06-95 11-08-95
WO 9411959 A	26-05-94	US 5345595 A AU 5596294 A CA 2149135 A EP 0669061 A JP 8503346 T US 5615408 A	06-09-94 08-06-94 26-05-94 30-08-95 09-04-96 25-03-97
US 5602906 A	11-02-97	NONE	

THIS PAGE BLANK (USPTO)